An Energy-Efficient VCO-Based Matrix Multiplier Block to Support On-Chip Image Analysis

Imon Banerjee and Arindam Sanyal

arXiv:1612.03361v1 [cs.ET] 11 Dec 2016

***Abstract*—Images typically are represented as uniformly sampled data in the form of matrix of pixels/voxels. Therefore, matrix multiply-and-accumulate (MAC) forms the core of most state-of-the-art image analysis algorithms. While digital implementation of MAC has generally been the preferred approach, high power consumption is an impediment to adopting it for medical image analysis. In this work, we present a time-domain signal processing architecture which performs MAC operations with 7bit accuracy while consuming 400X lower energy than digital implementation. The proposed architecture performs analog computation using mostly digital circuits and is suitable for scaled CMOS technologies. The proposed time-domain MAC architecture is expected to play a central role in empowering the advancement of various on-chip image analysis operations.**

***Index Terms*—On-chip image analysis, voltage-controlled oscillator, time-domain, matrix multiplication**

I. Introduction

In the digital data era, 2D/3D image analysis operations (e.g. registration, feature calculation, interpolation, fusion) are the core processing block of a wide range of automated systems, including computer aided diagnosis (CAD) [1], [2]. For example, in the modern CAD systems, an important image analysis operation is the co-registration of positron emission tomography (PET) image and magnetic resonance image (MRI) which combines functional information from PET images with anatomical information in MR images. The efficient co-registration of PET and MRI images can pave the way for a better understanding of physiological and disease mechanisms in pre-clinical and clinical settings.

The co-registration algorithm applies rigid registration [3] to map each voxel (x_1, y_1, z_1) in the MR image into the co-ordinate space of PET image (x_2, y_2, z_2), provided that the images of both modalities are acquired from the same subject and the scanning processes have not introduced nonrigid spatial transformations. To illustrate the algorithmic facets, we present a workflow diagram in Fig. 1 that takes MRI and PET images as inputs and uses rigid registration to compute the co-registered MRI/PET image. The rigid registration is done through multiply-and-accumulate (MAC) operations performed on the original voxel locations in 3D space (x,y,z) with 4×4 translation matrix (U) which represents translations(t) and scaling(S) and 4×4 rotation matrices (R) which represent rotation (see (1)).

I. Banerjee is with the Laboratory of Quantitative Imaging, Stanford University School of Medicine, Stanford 94305, CA, USA. (e-mail: imonb@stanford.edu)

A. Sanyal is with Department of Electrical Engineering, The State University of New York at Buffalo, Buffalo 14260, NY,USA. (email: arindams@buffalo.edu)

The transformation matrix (M = translation + rotation) for each voxel is derived automatically by analyzing the difference between source and target data of the registration, which can again be described as a sequence of MAC operations. Note that the MRI and the PET images with reasonable accuracy have number of voxels more than $512 \times 512 \times 50$. Therefore, a standard registration operation requires $> 13, 10, 7200$ MAC operations which can consume a significant amount of power and computation time. Similarly, other image analysis operations, such as object segmentation, feature extraction, can be decomposed into a set of sequential MAC operations. Hence, reduction in power consumption during MAC operations is a significant challenge that has to be addressed in order to develop on-chip image analysis blocks.

Spurred on by Moore's law and CMOS technology scaling, the general approach towards CAD has been to perform all mathematical operations digitally using computers. While digital computation can provide high accuracy, the power consumption has been prohibitively high to prevent portable solutions. With CMOS technology scaling slowing down, there is an interest in analog signal processing (ASP) to perform mathematical computations in a low power fashion [4]. A key enabler for ASP methodology is that most image processing algorithms require only 6-8 bits precision. ASP excels at approximate computing while consuming lower power than digital computing. Recent approaches to analog computation has been to use switches and capacitors as signal processing elements [5]–[8] in advanced CMOS technologies. The matrix multiplier reported in [8] performs calculations with 6b linearity while having an energy consumption of only 13fJ/operation at 1GHz. The work in [5] presents switched capacitor multipliers and dividers. However, it uses amplifiers which is not power efficient, specially in advanced CMOS technologies. The works in [6]–[8] present power efficient multipliers by doing away with amplifiers and relying on only switched capacitors for signal processing. These multipliers work in the charge domain and are good solutions for high speed operations usually in the range of GHz of bandwidths. However, charge leakage presents a significant challenge to using switched-capacitor multipliers for medical signal processing in which computations can require only a few MHz of bandwidths. Thus, switched capacitor multipliers are not very suitable for medical image processing.

In this letter, we present an alternate solution by performing multiplication and addition operations in time domain, which satisfies medical image processing requirements. The proposed technique has significantly higher immunity to charge leakage than switched capacitors and can trade-off speed for power

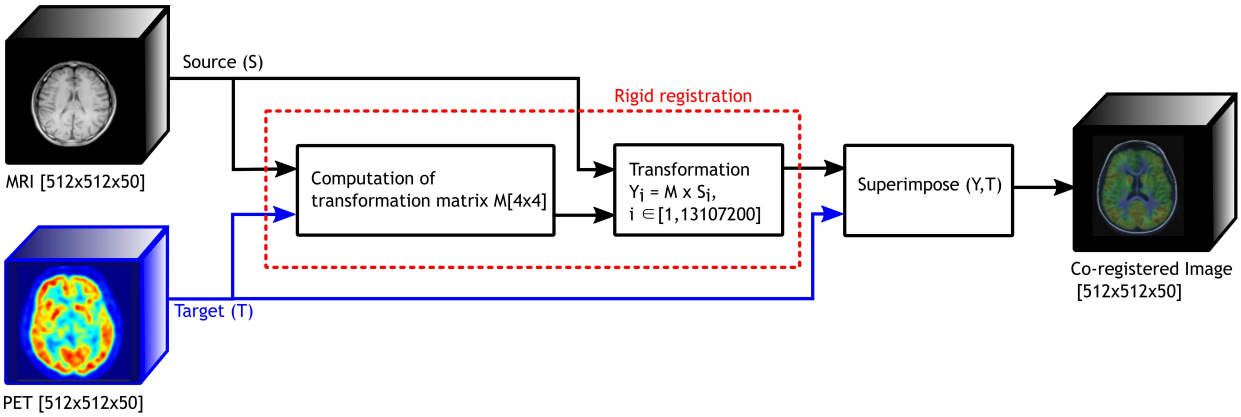

Fig. 1. A workflow diagram of MRI and PET brain image co-registration.

$$U = \begin{vmatrix} S_x & 0 & 0 & 0 \\ 0 & S_y & 0 & 0 \\ 0 & 0 & S_z & 0 \\ t_x & t_y & t_z & 1 \end{vmatrix}; R_x = \begin{vmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos\theta & -\sin\theta & 0 \\ 0 & \sin\theta & \cos\theta & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix}; R_y = \begin{vmatrix} \cos\theta & 0 & -\sin\theta & 0 \\ 0 & 1 & 0 & 0 \\ -\sin\theta & 0 & \cos\theta & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix}; R_x = \begin{vmatrix} \cos\theta & -\sin\theta & 0 & 0 \\ \sin\theta & \cos\theta & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix} \quad (1)$$

without sacrificing computational accuracy. Time domain circuits are uniquely suitable for scaled CMOS technology. They are highly digital and hence can operate at low supply voltages in advanced CMOS technologies. In addition, the quantization noise in time domain circuits is essentially transistor delay which reduces with technology scaling. Thus, the proposed architecture can be used for a power efficient hardware implementation of the rigid registration block (see Fig. 1) which is a core element in PET/MRI modality fusion process. In addition, the proposed time domain architecture holds the promise of providing low power computational ability for a wide range of portable CAD systems.

The rest of this letter is organized as follows: Section II presents a brief review of existing analog-to-digital matrix multipliers, Section III discusses the key idea behind the proposed time-domain matrix multiply-and-add operator, Section IV presents a CMOS circuit implementation and simulation results, while the conclusion and future research direction are brought up in Section V.

II. Review of Analog-to-Digital Matrix Multiplication and Addition

A MAC operation is defined as

$$Y = \sum_{j=1}^{N} W_j X_j \quad (2)$$

where X_j is the input and W_j is the weight.

Several analog signal processing techniques have been reported which accept an input X_j, perform the multiplication in (2) in analog domain and return a digital output $Y_j = W_j X_j$. A straightforward method is to use an operational amplifier to perform voltage domain multiplication and addition as described by (2). The operational amplifier needs to have a high gain to perform the MAC operation accurately. However, high gain operational amplifiers are power hungry and instead four-quadrant CMOS multipliers [9] are often used for approximate multiplication. To reach higher energy efficiency, [10] uses charge coupling to implement approximate analog matrix multipliers. [11] uses low performance, thin-film transistors (TFT) to implement an approximate analog multiplier for processing data from sensors.

More recently, there have been efforts to integrate analog-digital matrix multiplication (AD-MM) inside analog-to-digital converters (ADC) to lower power consumption. Capacitors and switches are used for passive multiplication in [7], [8], [12]. Charge domain passive multipliers achieve very low power consumption in advanced CMOS technology nodes as they do not use any active amplifiers which are power hungry. Addition can be done in current or charge domain in a low power fashion.

III. Time Domain Matrix Operations

While passive switched capacitor techniques are a good solution for low energy, approximate multipliers, they suffer from non-idealities due to charge leakage specially for advanced CMOS technologies which suffer from increased leakage. This problem is exacerbated at low speeds of operation, which implies that further lowering of energy consumption by reducing speed is challenging for these techniques. In addition, often switched capacitor multipliers are used in conjunction with voltage domain ADCs which suffer from reduced dynamic range as the supply voltage is scaled down in advanced CMOS technologies.

To counter these limitations of switched capacitor multipliers, we propose a time domain multiplier. By shifting to time domain, sensitivity to charge leakage is greatly diminished and very low energy multipliers can be designed which are

suitable for low bandwidth medical image processing. An added advantage is that quantization noise of time-domain multipliers come from transistor delay which reduces with technology scaling. Thus, in advanced CMOS technologies, time-domain multipliers are better suited for low-bandwidth operations than switched capacitors.

A voltage-controlled oscillator (VCO) is a major building block of time domain multiplier. The quantized phase of a VCO can be written as

$$\phi[k] = 2\pi K_v \int_0^{T_{int}} V_{in}(t)dt \tag{3}$$

where V_{in} is the input to the VCO, K_v is the VCO gain and T_{int} is time over which the VCO phase is integrated. If V_{in} changes slowly, $\phi[k]$ can be written as

$$\begin{aligned}\phi[k] &= 2\pi K_v \sum_{j=1}^{N}(t_j - t_{j-1})V_{in}[j] \\ &= 2\pi K_v \sum_{j=1}^{N}(\Delta t_j)\, V_{in}[j] \end{aligned} \tag{4}$$

where $T_{int} = \sum_{j=1}^{N} \Delta t_j$ and $V_{in}[j] = V_{in}(t = t_j)$

(4) is analogous to (2) with $W_j \equiv 2\pi K_v(\Delta t_j)$ and $X_j \equiv V_{in}[j]$. Thus, a VCO can be used to perform a MAC operation in phase domain. The equivalent digital output of the MAC operation can be readily obtained by simply sampling the output of the VCO, without requiring a separate ADC as in the charge-domain matrix multipliers. The accumulation operation is done in phase domain and is highly linear. Unlike charge-domain architectures, accumulation in phase domain comes without any additional power consumption. Nonlinearity in phase domain MAC operation is primarily due to nonlinearity in voltage-to-phase conversion. Increasing the integration time, T_{int}, allows reduction of VCO gain, K_v, which in turn increases VCO linearity. This is particularly suitable for medical image processing which does not require a high bandwidth, and hence, linearity of the VCO can be increased by lowering K_v and increasing T_{int}. As long as the VCO is oscillating, there is no leakage error in the phase value held by the VCO which is of importance for low bandwidth medical signal processing.

IV. Matrix Multiplier Architecture

Fig. 2 shows the conceptual block diagram of the proposed time-domain matrix multiplier architecture along with its timing diagram. A voltage-to-current (V/I) converter drives a current-controlled oscillator (CCO) and the quantized phase output of the CCO holds the result of multiplication of V_{in} and $2\pi K_v \phi_1$. The duration of ϕ_1 in the $j-$th sampling period is Δt_j which is digitally controllable. Addition comes without any hardware cost as the CCO holds on to its phase which keeps accumulating with time. Fig. 2 illustrates how the proposed architecture performs MAC operation.

Fig. 3 shows circuit implementation of the proposed architecture. The architecture is implemented in a differential

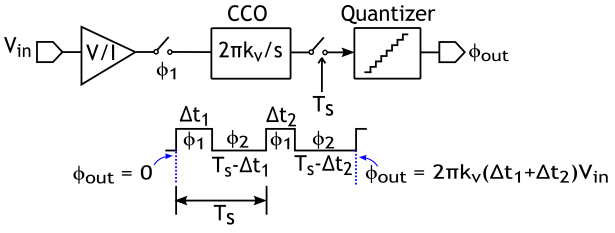

Fig. 2. Architecture and timing diagram of proposed time domain matrix multiplier.

fashion to suppress common-mode noise on the inputs as well as noise from supply and ground. The input signal is applied differentially to a V/I converter. The V/I converter drives two pseudo-differential CCOs during the phase ϕ_1 and the two CCOs are run with a low current supply I_L during the phase ϕ_2. The two CCOs are not stopped during ϕ_2 to ensure that the accumulated phase held by them are not corrupted by leakage. By running the two CCOs at the same frequency during ϕ_2, no phase is accumulated differentially at ϕ_2. The output of each CCO stage is latched by the sampling clock into a flip-flop (FF). At any given time, only one of the CCO stages is in a state of either a positive or negative transition. Thus, for an N-stage CCO, the instantaneous phase can be quantized into $2N$ levels between $(0, 2\pi)$ corresponding to N positive transitions and N negative transitions. The matrix multiplier is highly digital and makes use of simple digital circuits to perform time domain analog signal processing.

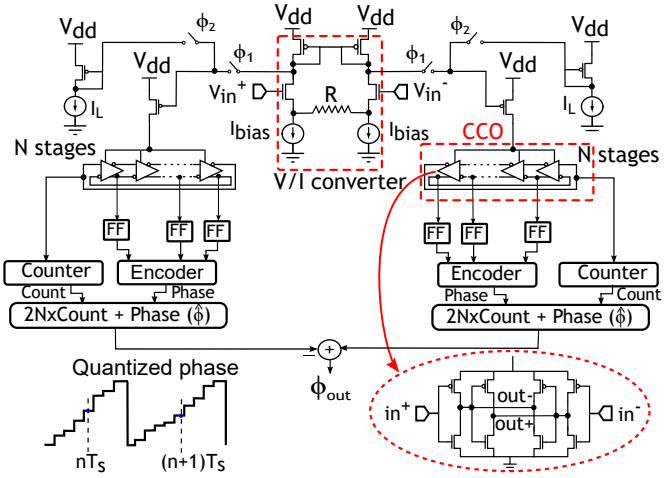

Fig. 3. Circuit schematic of proposed time domain matrix multiplier.

As shown in Fig. 3, CCO phase increases monotonically with time and wraps over after it crosses 2π. A counter is used to store the number of times the VCO phase overflows over the period of integration. The total phase at any time is given by $(2N \cdot Count + \hat{\phi})$ where $\hat{\phi}$ is the instantaneous quantized phase.

Since VCO gain varies with process, voltage and temperature (PVT), the result of multiply-and-add will vary with PVT. Hence, background tracking is needed to correct for VCO gain variation. The V/I transconductance is given by $1/R$ for $g_m R \gg 1$, where g_m is the transconductance of the V/I input transistors and R is their source degeneration resistance. Thus,

the V/I converter is relatively insensitive to PVT variations, and hence, resistance trimming is not required for multipliers with 6-8 bits accuracy. To track the CCO gain, width of the tail current source can be changed depending on the output of a counter which is clocked by the CCO as shown in Fig. 4. The counter output is monitored by a comparator which is clocked by a divided-down version of the sampling clock. The counter is reset after every comparison. If the CCO is running too fast, the comparator will reduce the tail current source width and vice versa. This ensures that the counter output is held equal to a preset value F_{in} which sets the CCO free-running frequency.

The background tracking technique can be applied to a reference multiplier and the comparator digital output word can be applied to the tail current source of all the CCOs. Process related mismatch between the tail current sources of the different CCOs will limit the accuracy to which the PVT sensitivity can be corrected. Fortunately, the matching accuracy is not very stringent as only 6-8 bits accuracy is required. In addition, CCO tail current devices are usually made large to reduce flicker noise. Large device size also reduces mismatches. For a large design with many multiplier cells, a few local copies of the reference multiplier can be distributed across the chip to account for gradient mismatches.

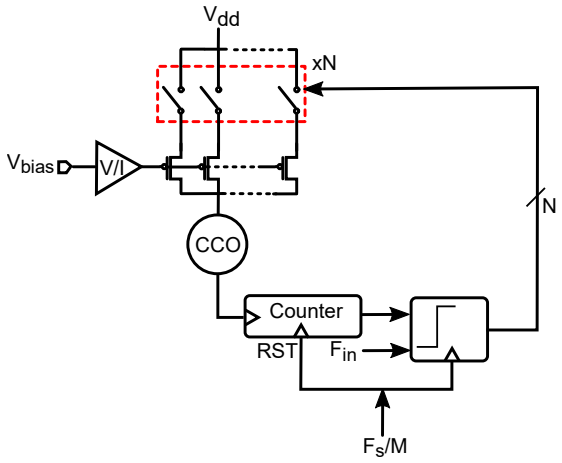

Fig. 4. VCO gain tracking architecture.

A MAC cell was designed in 40nm CMOS process. With a 400mV differential input running at 0.6MHz, the MAC cell operates with 7-bit linearity. To test the accuracy of the proposed MAC cell, a row vector (1×512) is multiplied with a column vector (512×1). The row vector $\begin{bmatrix} x_0 & x_1 & x_2 & \cdots & x_{511} \end{bmatrix}$ is set to $\begin{bmatrix} 0 & \sin(\omega T_s) & \sin(2\omega T_s) & \cdots & \sin(511\omega T_s) \end{bmatrix}$ and all the elements in the column vector are set to T_s where T_s = 10ns. Fig. 5 shows the simulated result for multiplication of the row vector and the column vector as well as the error between the output of the MAC operation and the desired output. It can be seen from Fig. 5 that the proposed MAC cell has a low quantization error and has 7 bit linearity.

The MAC cell consumes 86μW from a 1.1V supply while operating at 100 MHz. The energy efficiency of the proposed time-domain MAC cell is 2fJ/operation, compared to 968 fJ/operation energy efficiency of highly optimized digital matrix multipliers [13]. Thus, the proposed MAC cell has more than 400X better energy efficiency than digital matrix multipliers.

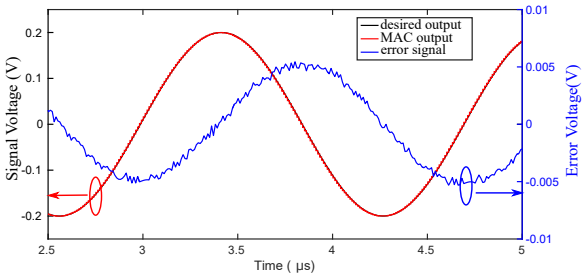

Fig. 5. Matrix multiply-and-add transient simulation.

V. Conclusion

A time-domain architecture for performing multiplication-and-addition operations is presented in this letter. The proposed architecture exploits the low bandwidth and not-too stringent accuracy requirement of medical image processing algorithms to achieve drastic increase in energy efficiency compared to digital matrix multipliers. The architecture is highly digital and suitable for advanced CMOS technologies. The proposed architecture can act as an enabler for developing portable hardware solutions for 2D/3D image analysis.

References

[1] Banerjee, I., Agibetov, A., Catalano, C.E., Patanè, G., Spagnuolo, M., 'Semantics-driven Annotation of Patient-Specific 3D Data: A Step to Assist Diagnosis and Treatment of Rheumatoid Arthritis', *The Visual Computer Journal*, 2016, pp. 1–13.

[2] Slomka, P. J.: 'Software approach to merging molecular with anatomic information', *Journal of Nuclear Medicine*, 2004, vol. 45, pp. 36S–45S.

[3] Goshtasby, A. Ardeshir.: '2-D and 3-D image registration: for medical, remote sensing, and industrial applications', *John Wiley & Sons*, 2005.

[4] St. Amant, R., Yazdanbakhsh, A., Park, J., Thwaites, B., Esmaeilzadeh, H., Hassibi, A., Ceze, L., and Burger, D.: 'General-purpose code acceleration with limited-precision analog computation', *ACM SIGARCH Comput. Archit. News* , 2014, vol. 42, no. 3, pp. 505–516.

[5] Watanabe, K., and Temes, G.: 'A switched-capacitor multiplier/divider with digital and analog output', *IEEE Trans. Circuits Syst.*, 1984, vol. 31, no. 9, pp. 796–800.

[6] Sadhu, B., Sturm, M., Sadler, B.M., and Harjani, R.: 'Analysis and design of a 5GS/s analog charge-domain FFT for an SDR front-end in 65nm CMOS', *IEEE Journal of Solid-State Circuits*, 2013, vol. 48, no. 5, pp. 1199–1211.

[7] Bankman, D. and Murmann, B.: 'Passive charge redistribution digital-to-analogue multiplier', *Elec. Lett.*, 2015, vol. 51, no. 5, pp. 386–388.

[8] Lee, E.H., and Wong, S.S.: 'A 2.5GHz 7.7TOPS/W switched-capacitor matrix multiplier with co-designed local memory in 40nm', *IEEE Journal of Solid-State Circuits*, 2016, pp. 418–420.

[9] Bult, K., and Wallinga, H.: 'A CMOS four-quadrant analog multiplier', *IEEE Journal of Solid-State Circuits*, 1986, SSC-21, no. 3, pp. 430–445.

[10] Genov, R., and Cauwenberghs, G.: 'Kerneltron: support vector machine in silicon', *IEEE Trans. Neural. Netw.*, 2003, vol. 14, no. 5, pp. 1426–1434.

[11] Rieutort-Louis, W.R., Moy, T., Wang, Z., Wagner, S., Sturm, J.C., and Verma, N.: 'A large-area image sensing and detection system based on embedded thin-film classifiers', *IEEE Journal of Solid-State Circuits*, 2016, vol. 51, no. 1, pp. 281–290.

[12] Wang, Z., Zhang, J. and Verma,N.: 'Realizing low-energy classification systems by implementing matrix multiplication directly within an ADC', *IEEE Trans. Biomed. Circ. and Syst.* , 2015, vol. 9, no. 6, pp. 825–837.

[13] Saha, P., Banerjee, A., Bhattacharyya, P., and Dandapat, A.: 'Improved matrix multiplier design for high-speed digital signal processing applications', *IET Circuits, Devices and Systems*, 2014, vol. 8, no. 1, pp. 27–37.